616.301 FC Advanced Microeconomics

John Hillas

The University of Auckland

Department of Economics

First Semester, 1999

VISIT…

Contents

Chapter 1. General Equilibrium Theory 3
1. The basic model of a competitive economy 4
2. Walrasian equilibrium 8
3. Edgeworth Boxes 10
4. The First and Second Fundamental Theorems of Welfare Economics 17
5. Exercises 20

Chapter 2. Noncooperative Game Theory 28
1. Normal Form Games 30
2. Extensive Form Games 42
3. Existence of Equilibrium 54

Chapter 3. Auctions 59
1. Introduction 60
2. Types of Auction 62
3. Analysis of the Auctions I: Private Values 66
4. Bibliography 78

CHAPTER 1

General Equilibrium Theory

There are two central multi-agent models used in economics: the general equilibrium model and the strategic or game theoretic model. In the strategic model we say what each actor in the economy (or in the part of the economy under consideration) can do. Each agent acts taking into consideration the plans of each other agent in the economy. There is a certain coherence to this. It is clearly specified what each person knows and how knowledge flows from one to another. It becomes difficult to specify in a completely satisfying way all the relevant details of the economy.

In the general equilibrium model on the other hand each actor does not take explicitly into account the actions of each other. Rather we assume that each reacts optimally to a market aggregate, the price vector. In comparison with the strategic model there is a certain lack of coherence. It is not specified exactly how the consumers interact or how information flows from one consumer to another. On the other hand the very lack of detail can be seen as a strength of the model. Since the details of how the actors interact is not specified we do not get bogged down in the somewhat unnatural details of a particular mode of interaction.

1. The basic model of a competitive economy

We summarise the basic ingredients of the model. All the following items except the last are part of the exogenous description of the economy. The price vector is endogenous. That is it will be specified as part of the *solution* of the model.

- L goods
- N consumers — a typical consumer is indexed consumer n. The set of all consumers is (abusively) denoted N. (That is, the same symbol N stands for *both* the number of consumers and the set of all consumers. This will typically not cause any confusion and is such common practice that you should become used to it.)
- the consumption set for each consumer is $\mathbb{R}^L_+$, the set of all L-dimensional vectors of nonnegative real numbers.
- $\succsim_n$ the rational preference relation of consumer n on $\mathbb{R}^L_+$ or u_n a utility function for consumer n mapping $\mathbb{R}^L_+$ to $\mathbb{R}$ the set of real numbers. That is, for any consumption bundle $x = (x_1, \ldots, x_L) \in \mathbb{R}^L_+$ u_n tells us the utility that consumer n associates to that bundle.
- $\omega_n = (\omega_{1n}, \omega_{2n}, \ldots, \omega_{Ln})$ in $\mathbb{R}^L_+$ the endowment of consumer n
- p in $\mathbb{R}^L_{++}$ a strictly positive price vector; $p = (p_1, \ldots, p_\ell, \ldots, p_L)$ where $p_\ell > 0$ is the price of the ℓth good.

Definition 1.1. An *allocation*

$$x = ((x_{11}, x_{21}, \ldots, x_{L1}), \ldots, (x_{1N}, x_{2N}, \ldots, x_{LN}))$$

in $(\mathbb{R}^L_+)^N$ specifies a consumption bundle for each consumer. A *feasible allocation* is an allocation such that

$$\sum_{n \in N} x_n \leq \sum_{n \in N} \omega_n$$

or equivalently that, for each ℓ

$$\sum_{n \in N} x_{\ell n} \leq \sum_{n \in N} \omega_{\ell n}.$$

That is, for each good, the amount that the consumers together consume is no more than the amount that together they have. (Note that we are implicitly assuming that the goods are freely disposable. That is, we do not assume that all the good is necessarily consumed. If there is some left over it is costlessly disposed of.)

Definition 1.2. Consumer n's *budget set* is

$$B(p, \omega_n) = \{x \in \mathbb{R}^L_+ \mid p \cdot x \leq p \cdot \omega_n\}$$

Thus the budget set tells us all the consumption bundles that the consumer could afford to buy at prices $p = (p_1, p_2, \ldots, p_\ell, \ldots, p_L)$ if she first sold all of her endowment at those prices and funded her purchases with the receipts. Since we assume that the consumer faces the same prices when she sells as when she buys it does not make any difference whether we think of her as first selling all of her endowment and then buying what she wants or selling only part of what she has and buying a different incremental bundle to adjust her overall consumption bundle.

DEFINITION 1.3. Consumer n's *demand correspondence* is

$$x_n(p, \omega_n) = \{x \in B(p, \omega_n) \mid \text{there is no } y \in B(p, \omega_n) \text{ with } y \succ_n x\}$$

or, in terms of the utility function

$$x_n(p, \omega_n) = \{x \in B(p, \omega_n) \mid \text{there is no } y \in B(p, \omega_n) \text{ with } u_n(y) > u_n(x)\}.$$

In words we say that the demand correspondence for consumer n is a rule that associates to any price vector the set of all affordable consumption bundles for consumer n for which there is no affordable consumption bundle that consumer n would rather have.

Let us now make some fairly strong assumptions about the $\succsim_n$'s, or equivalently, the utility functions u_n. For the most part the full strength of these assumptions is unnecessary. Most of the results that we give are

true with weaker assumptions. However these assumptions will imply that the demand correspondences are, in fact, functions, which will somewhat simplify the presentation.

We assume that for each n the preference relation $\succsim_n$ is (a) continuous (this is technical and we won't say anything further about it), (b) strictly increasing (if $x \geq y$ and $x \neq y$ then $x \succ_n y$), and (c) strictly convex (if $x \succsim_n y$, $x \neq y$, and $\alpha \in (0,1)$ then $\alpha x + (1-\alpha)y \succ_n y$).

If we speak instead of the utility functions then we assume that the utility function u_n is (a) continuous (this is again technical, but you should know what a continuous function is), (b) strictly increasing (if $x \geq y$ and $x \neq y$ then $u_n(x) > u_n(y)$), and (c) strictly quasi-convex (if $u_n(x) \geq u_n(y)$, $x \neq y$, and $\alpha \in (0,1)$ then $u_n(\alpha x + (1-\alpha)y) > u_n(y)$).

PROPOSITION 1.1. *If $\succsim_n$ is continuous, strictly increasing, and strictly convex (u_n is continuous, strictly increasing, and strictly quasi-convex) then*

1. *$x_n(p, \omega_n) \neq \emptyset$ for any ω_n in $\mathbb{R}^L_+$ and any p in $\mathbb{R}^L_{++}$,*
2. *$x_n(p, \omega_n)$ is a singleton so $x_n(\cdot, \omega_n)$ is a function, and*
3. *$x_n(\cdot, \omega_n)$ is a continuous function.*

2. Walrasian equilibrium

We come now to the central solution concept of general equilibrium theory, the concept of competitive or *Walrasian equilibrium*. Very briefly a Walrasian equilibrium is a situation in which total demand does not exceed total supply. Indeed, is all goods are desired in the economy, as we assume they are, then it is a situation in which total demand exactly equals total supply. We state this more formally in the following definition.

DEFINITION 1.4. The price vector p is a *Walrasian (or competitive) equilibrium price* if

$$\sum_{n\in N} x_n(p,\omega_n) \leq \sum_{n\in N} \omega_n.$$

If we do *not* assume that the demand functions are single valued the we need a slightly more general form of the definition.

DEFINITION 1.4′. The pair (p,x) in $\mathbb{R}^L_{++} \times (\mathbb{R}^L_+)^N$ is a *Walrasian equilibrium* if x is a feasible allocation (that is, $\sum_{n\in N} x_n \leq \sum_{n\in N} \omega_n$) and, for each n in N

$$x_n \succsim_n y \text{ for all } y \text{ in } B(p,\omega_n).$$

Since we assume that $\succsim_n$ is strictly increasing (in fact local nonsatiation is enough) it is fairly easy to see that the only feasible allocations that will be involved in any equilibria are those for which

$$\sum_{n\in N} x_n = \sum_{n\in N} \omega_n. \tag{1}$$

3. Edgeworth Boxes

We shall now examine graphically the case $L = N = 2$. An allocation in this case is a vector in $\mathbb{R}^4_+$. However, since we have the two equations of the vector equation 1 we can eliminate two of the variables and illustrate the allocations in two dimensions. A particularly meaningful way of doing this is by what is known as the Edgeworth box.

Let us first draw the consumption set and the budget set for each consumer, as we usually do for the two good case in consumer theory. We show this in Figure 1 and Figure 2. The only new feature of this graph is that rather than having a fixed amount of wealth each consumer starts off with an initial endowment bundle ω_n. The boundary of their budget set (that is, the budget line) is then given by a line through ω_n perpendicular to the price vector p.

What we want to do is to draw Figure 1 and Figure 2 in the same diagram. We do this by rotating Figure 2 through 180° and then lining the figures up so that ω_1 and ω_2 coincide. We do this in Figure 3. Any point x in the diagram now represents (x_{11}, x_{21}) if viewed from 0_1 looking up with the normal perspective *and simultaneously* represents (x_{12}, x_{22}) if viewed from 0_2 looking down. Notice that while all the feasible allocations are within the "box" part of each consumer's budget set goes outside the "box." One of the central ideas of general equilibrium theory is that the decision making can be decentralized by the price mechanism. Thus neither consumer is required to take into account when making their choices what

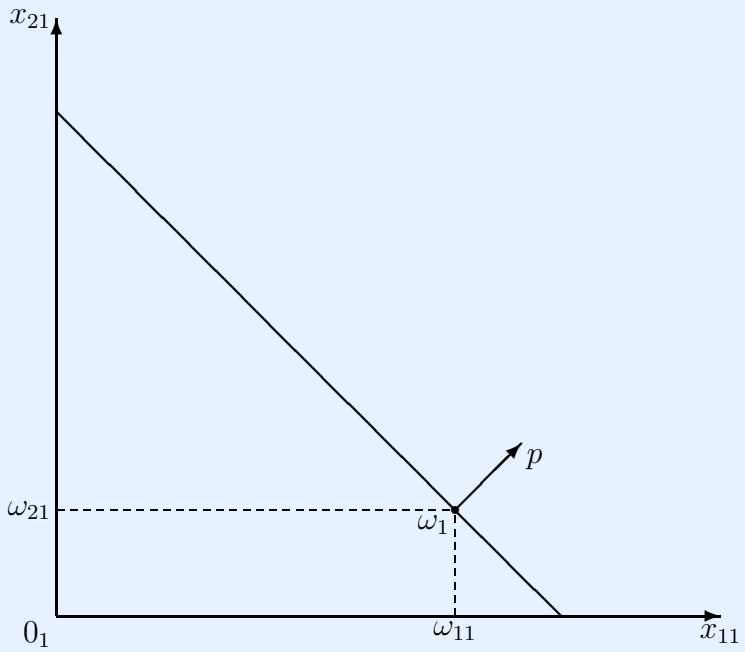

Figure 1

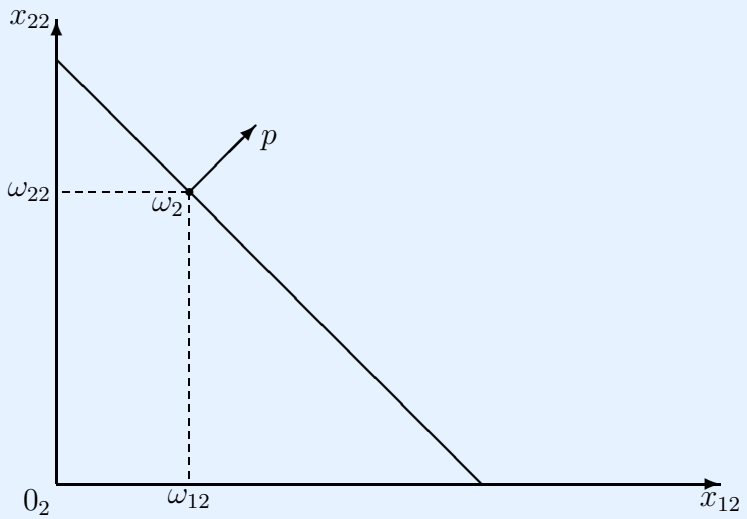

Figure 2

is globally feasible for the economy. Thus we really do want to draw the diagrams as I have and not leave out the parts "outside the box."

We can represent preferences in the usual manner by indifference curves. I shall not again draw separate pictures for consumers 1 and 2, but rather go straight to drawing them in the Edgeworth box, as in Figure 4.

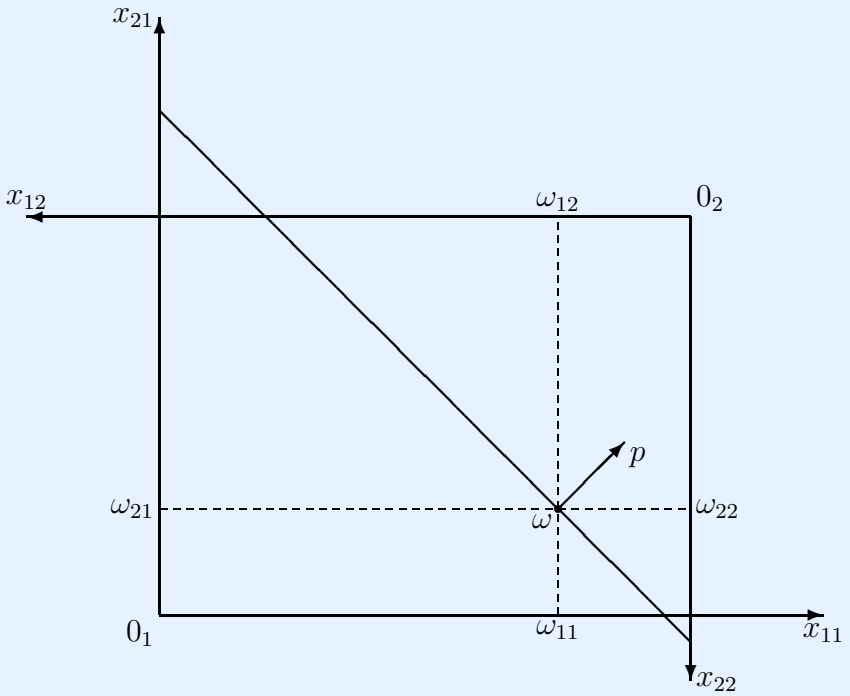

Figure 3

Let us look at the definition of a Walrasian equilibrium. If some allocation feasible x ($\neq \omega$) is to be an equilibrium allocation then it must be

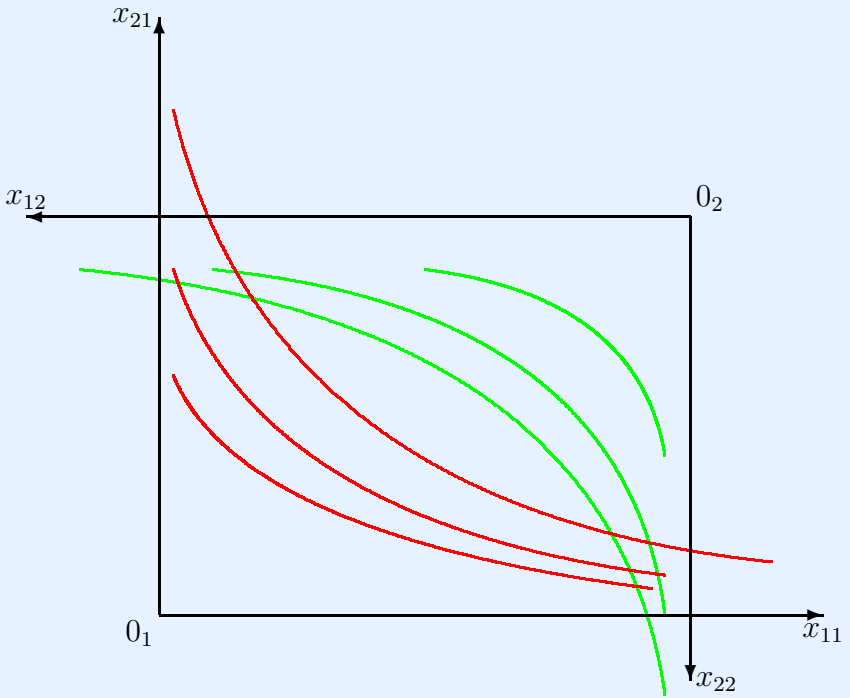

Figure 4

in the budget sets of both consumers. (Such an allocation is shown in Figure 5.) Thus the boundary of the budget sets must be the line through x and ω (and the equilibrium price vector will be perpendicular to this line).

Also x must be, for each consumer, at least as good as any other bundle in their budget set. Now any feasible allocation y that makes Consumer 1 better off than he is at allocation x must not be in Consumer 1's budget set. (Otherwise he would have chosen it.) Thus the allocation must be strictly above the budget line through ω and x. But then there are points in Consumer 2's budget set which give her strictly more of both goods than she gets in the allocation y. So, since her preferences are strictly increasing there is a point in her budget set that she strictly prefers to what she gets in the allocation y. But since the allocation x is a competitive equilibrium with the given budget sets then what she gets in the allocation x must be at least as good any other point in her budget set, and thus strictly better than what she gets at y.

What have we shown? We have shown that if x is a competitive allocation from the endowments ω then any feasible allocation that makes Consumer 1 better off makes Consumer 2 worse off. We can similarly show that any feasible allocation that makes Consumer 2 better off makes Consumer 1 worse off. In other words x is Pareto optimal.

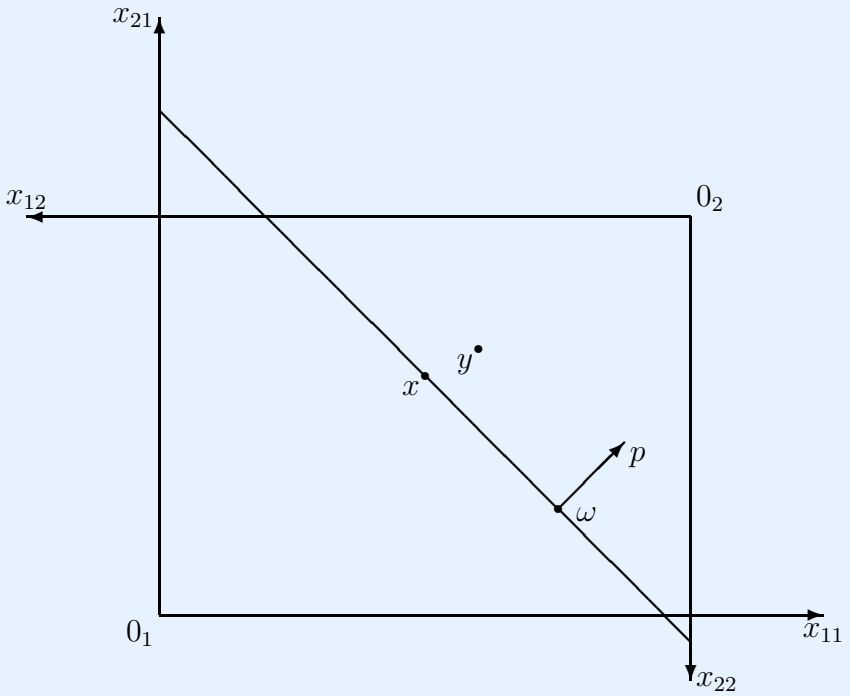

Figure 5

4. The First and Second Fundamental Theorems of Welfare Economics

We shall now generalise this intuition into the relationship between equilibrium and efficiency to the more general model. We first define more formally our idea of efficiency.

DEFINITION 1.5. A feasible allocation x is *Pareto optimal* (or *Pareto efficient*) if there is no other feasible allocation y such that $y_n \succsim_n x_n$ for all n in N and $y_{n'} \succ_{n'} x_{n'}$ for at least one n' in N.

In words we say that a feasible allocation is Pareto optimal if there is no other feasible allocation that makes at least one consumer strictly better off without making any consumer worse off. The following result generalises our observation about the Edgeworth box.

THEOREM 1.1 (The First Fundamental Theorem of Welfare Economics). *Suppose that for each n the preferences $\succsim_n$ are strictly increasing and that (p, x) is a Walrasian equilibrium. Then x is Pareto optimal.*

In fact, we can say something in the other direction as well. It clearly is not the case that any Pareto optimal allocation is a Walrasian equilibrium. A Pareto optimal allocation may well redistribute the goods, giving more to some consumers and less to others. However, if we are permitted to make such transfers then any Pareto optimal allocation is a Walrasian equilibrium from some redistributed initial endowment. Suppose that in

the Edgeworth box there is some point such as x in Figure 6 that is Pareto optimal. Since x is Pareto optimal Consumer 2's indifference curve through x must lie everywhere below Consumer 1's indifference curve through x. Thus the indifference curves must be tangent to each other. Let's draw the common tangent. Now, if we redistribute the initial endowments to some point ω' on this tangent line then with the new endowments the allocation x is a competitive equilibrium. This result is true with some generality, as the following result states. However we do require stronger assumptions that were required for the first welfare theorem. We shall look below at a couple of examples to illustrate why these stronger assumptions are needed.

THEOREM 1.2 (The Second Fundamental Theorem of Welfare Economics). *Suppose that for each n the preferences $\succsim_n$ are strictly increasing, convex, and continuous and that x is Pareto optimal with $x > 0$ (that is $x_{\ell n} > 0$ for each ℓ and each n. Then there is some feasible reallocation ω' of the endowments (that is $\sum_{n\in N} \omega'_n = \sum_{n\in N} \omega_n$) and a price vector p such that (p, x) is a Walrasian equilibrium of the economy with preferences $\succsim_n$ and initial endowments ω'.*

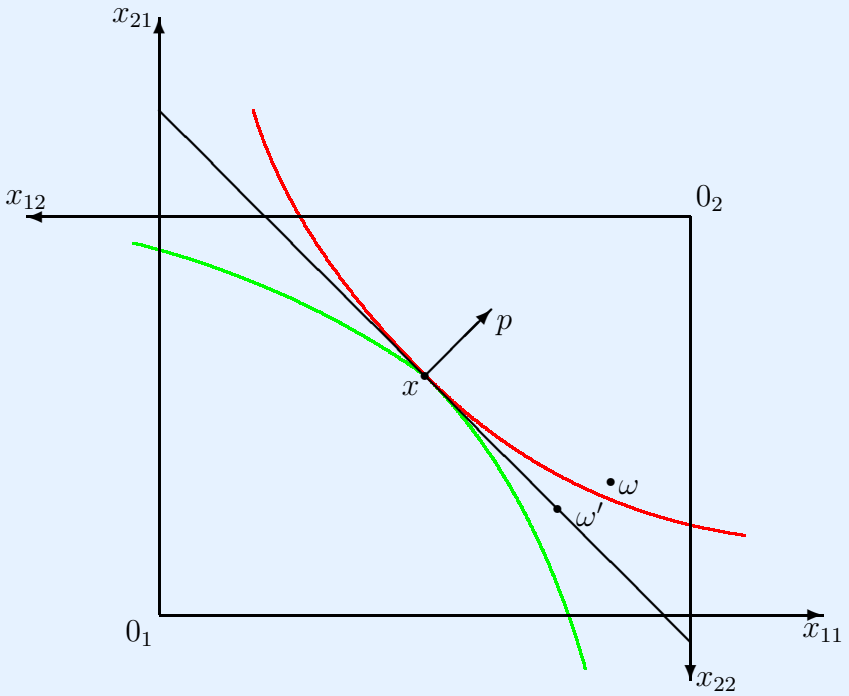

Figure 6

5. Exercises

EXERCISE 1.1. Consider a situation in which a consumer consumes only two goods, good 1 and good 2, and takes prices as exogenously given. Suppose that the consumer is initially endowed with ω_1 units of good 1 and ω_2 units of good 2, and that the price of good 1 is p_1 and the price of good 2 is p_2. (The consumer can either buy or sell at these prices.)

Suppose that the consumer's preferences are given by the utility function

$$u(x_1, x_2) = x_1 x_2$$

where x_1 is the amount of good 1 the consumer consumes and x_2 the amount of good 2. Suppose also that $\omega_1 = 2$ and $\omega_2 = 1$ and that $p_2 = 1$.

1. Write the budget constraint of this consumer and solve the budget constraint to give x_2 as a function of x_1.
2. Substitute this function into the utility function to give utility as a function of x_1.
3. Find the value of x_1 that maximises the consumer's utility. You do this by differentiating the function you found in the previous part setting the derivative equal to zero and then solving the resulting equation to give x_1 as a function of p_1.
4. Substitute the value you find for x_1 into the function you found in part 1 to find the optimal value of x_2 as a function of p_1.
5. Graph the functions $x_1(p_1)$ and $x_2(p_1)$.

6. Find the value of p_1 that will make this consumer willing to consume her initial endowment.
7. Suppose that p_2 was 2, rather than 1. Repeat the analysis above for this case. Comment on the result.

EXERCISE 1.2. Suppose that in addition to the consumer described in the previous exercise we also have an another consumer, consumer 2, whose preferences are given by the utility function

$$u_2(x_1, x_2) = {x_1}^2 x_2$$

where x_1 is the amount of good 1 the consumer consumes and x_2 the amount of good 2. Suppose also that $\omega_{12} = 1$ and $\omega_{22} = 3$, where $\omega_{\ell 2}$ is consumer 2's initial endowment of good ℓ. Again assume that $p_2 = 1$.

1. Repeat the analysis of the previous exercise for this consumer, finding the demand functions $x_{12}(p_1)$ and $x_{22}(p_1)$.
2. Find the market demand in the economy consisting of these two consumers by adding the individual demand functions. That is

   $$X_1(p_1) = x_{11}(p_1) + x_{12}(p_1)$$

   and

   $$X_2(p_1) = x_{21}(p_1) + x_{22}(p_1).$$

3. Find the value of p_1 for which the market demand for good 1 is exactly equal to the total endowment of good 1.

4. Confirm that at this price the the market demand for good 2 is also equal to the total endowment of good 2.
5. Why?

EXERCISE 1.3. Consider a situation in which a consumer consumes only two goods, good 1 and good 2, and takes prices as exogenously given. Suppose that the consumer is initially endowed with ω_1 units of good 1 and ω_2 units of good 2, and that the price of good 1 is p_1 and the price of good 2 is p_2. (The consumer can either buy or sell at these prices.)

Suppose that the consumer's preferences are given by the utility function

$$u(x_1, x_2) = \max\{x_1, x_2\}$$

where x_1 is the amount of good 1 the consumer consumes and x_2 the amount of good 2. [If you have come across Leontief preferences in the past be careful. These are almost precisely the opposite. Leontief preferences are defined by $u(x_1, x_2) = \min\{x_1, x_2\}$. You are asked to analyse Leontief preferences in Exercise 1.5.] Suppose also that $\omega_1 = 2$ and $\omega_2 = 3$. Again, we can choose one normalisation for the price vector and again we choose to let $p_2 = 1$.

1. On a graph draw several indifference curves of this consumer.
2. Draw, on the same graph draw the budget sets when $p_1 = 0.5$, when $p_1 = 1$, and when $p_1 = 2$,

3. Find the value(s) of (x_1, x_2) that maximises the consumer's utility, for each of these budget sets. You don't need to do any differentiation to do this. Simply thinking clearly should tell you the answer.
4. Generalise the previous answer to give the optimal values of x_1 and x_2 as functions of p_1.
5. Graph the functions $x_1(p_1)$ and $x_2(p_1)$.
6. Argue that there is no value of p_1 that will make the consumer willing to consume her initial endowment, and that thus, in the one consumer economy there is no equilibrium.

EXERCISE 1.4. Consider again the preferences given in Exercise 1.3. We shall now illustrate the convexifying effect of having many consumers.

1. Suppose that there are two consumers identical to the consumer described in Exercise 1, that is they have the same preferences and the same initial endowment. Graph the aggregate demand functions

$$X_1(p_1) = x_{11}(p_1) + x_{12}(p_1)$$

and

$$X_2(p_1) = x_{21}(p_1) + x_{22}(p_1)$$

taking particular care at the value $p_1 = 1$.
2. Is there a competitive equilibrium for this economy? Why?
3. How many identical consumers of this type are needed in order that the economy should have an equilibrium?

EXERCISE 1.5. Consider a situation in which a consumer consumes only two goods, good 1 and good 2, and takes prices as exogenously given. Suppose that the consumer is initially endowed with ω_1 units of good 1 and ω_2 units of good 2, and that the price of good 1 is p_1 and the price of good 2 is p_2. (The consumer can either buy or sell at these prices.)

Suppose that the consumer's preferences are given by the utility function

$$u(x_1, x_2) = \min\{x_1, x_2\}$$

where x_1 is the amount of good 1 the consumer consumes and x_2 the amount of good 2. Think a little about what this utility function means. If the consumer has 3 units of good 1 and 7 units of good 2 then her utility is 3. If she has 3 units of good 1 and 17 units of good 2 then her utility is still 3. Increasing the amount of the good she has more of does not increase her utility. If she has 4 units of good 1 and 7 units of good 2 then her utility is 4. Increasing the amount of the good she has less of does increase her utility.

1. On a graph draw several indifference curves of this consumer.
2. Now assume that $p_2 = 1$ and $\omega = (1, 2)$. Draw, on the same graph draw the budget sets when $p_1 = 0.5$, when $p_1 = 1$, and when $p_1 = 2$,
3. Find the value(s) of (x_1, x_2) that maximises the consumer's utility, for each of these budget sets. You don't need to do any differentiation to do this. Simply thinking clearly should tell you the answer.

4. Generalise the previous answer to give the optimal values of x_1 and x_2 as functions of p_1.
5. Again generalise your answer by dropping the assumptions that $p_2 = 1$ and that $\omega = (1, 2)$ to give the optimal values of x_1 and x_2 as functions of p_1, p_2, ω_1, and ω_2.

EXERCISE 1.6. Suppose that in addition to the consumer described in Exercise 1.5 we also have an another consumer, consumer 2, whose preferences are given by the utility function

$$u_2(x_1, x_2) = x_1 x_2$$

where x_1 is the amount of good 1 the consumer consumes and x_2 the amount of good 2. Let $\omega_{\ell 2}$ be consumer 2's initial endowment of good ℓ. [This is the utility function you analysed in Homework 1.] Similarly let $u_1(x_1, x_2) = \min\{x_1, x_2\}$ be the utility function for consumer 1 that you discussed in the previous exercise and $\omega_{\ell 1}$ be consumer 1's initial endowment of good ℓ.

1. Repeat the analysis of the previous exercise for consumer 2, finding the demand functions $x_{12}(p_1, p_2, \omega_{12}, \omega_{22})$ and $x_{22}(p_1, p_2, \omega_{12}, \omega_{22})$.
2. Find the market demand in the economy consisting of these two consumers by adding the individual demand functions. That is

$$X_1(p_1, p_2, \omega_{11}, \omega_{21}, \omega_{12}, \omega_{22}) = x_{11}(p_1, p_2, \omega_{11}, \omega_{21}) + x_{12}(p_1, p_2, \omega_{12}, \omega_{22})$$

and

$$X_2(p_1, p_2, \omega_{11}, \omega_{21}, \omega_{12}, \omega_{22}) = x_{21}(p_1, p_2, \omega_{11}, \omega_{21}) + x_{22}(p_1, p_2, \omega_{12}, \omega_{22}).$$

3. Suppose that $\omega_{11} = 0$, $\omega_{21} = 1$, $\omega_{12} = 1$, and $\omega_{22} = 0$. Also normalise prices so that $p_2 = 1$. Find the value of p_1 for which the market demand for good 1 is exactly equal to the total endowment of good 1.
4. Confirm that at this price the the market demand for good 2 is also equal to the total endowment of good 2.
5. What is the utility of consumer 1 in this equilibrium?

EXERCISE 1.7. Suppose that everything is as in Exercise 1.6 except that now the initial endowments are $\omega_{11} = 0$, $\omega_{21} = 2$, $\omega_{12} = 1$, and $\omega_{22} = 0$. That is consumer 1's endowment of good 2 has increased from 1 to 2.

1. Find the equilibrium in this case.
2. What is the utility of consumer 1 in this equilibrium?
3. Compare this with the utility found in the previous exercise and explain why we get this apparently paradoxical result.
4. Explain why it should be possible to come up with an even stronger paradox, that is why we might expect to find examples in which the weak inequality you observed in the previous part could be strong. [You might do this by drawing the appropriate Edgeworth box.]

EXERCISE 1.8. Suppose that everything is as in Exercises 1.6 and 1.7 except that now the initial endowments are $\omega_{11} = 0$, $\omega_{21} = 0.25$, $\omega_{12} = 1$, and $\omega_{22} = 0$.

1. Is there an equilibrium in this case? [Hint: There is. One of the prices is zero.]
2. Let the initial endowments be $\omega_{11} = 0$, $\omega_{21} = a$, $\omega_{12} = 1$, and $\omega_{22} = 0$. Graph the equilibrium value of p_1 (normalising p_2 to be 1) and x_{11} the equilibrium consumption of good 1 by consumer 1 as functions (or correspondences) of a.

CHAPTER 2

Noncooperative Game Theory

As in the previous chapter we continue to examine situations in which a number of decision makers interact. There are a number of basic ways of modelling such situations. In some situations it is reasonable to assume that each decision maker reacts to not his assessment of what each individual will do but rather the value of some aggregate statistic which varies little with the choices of one individual. In such a case it is a reasonable modelling strategy to model the decision makers as taking as given the value of the aggregate variable. The main approach of this kind is general equilibrium theory which we examined in Chapter 1.

This chapter examines another method of modelling situations in which a number of decision makers interact. We shall think of the decision makers as acting based on their assessment of what each other *individual* decision maker will do. This approach is known as game theory or occasionally, and more informatively, as interactive decision theory.

There are two central models of interactive decision problems or games: the model of normal form games, and the model of extensive form games. While there exist minor variants we shall define in this section two rather standard versions of the models. We shall assume that the number of

decision makers or players is finite and that the number of choices facing each player is also finite.

1. Normal Form Games

1.1. Definition. Let's first give the definition of a finite normal form game and then discuss what each of the various parts means.

DEFINITION 2.1. A (finite) *normal form game* is a triple (N, S, u) where $N = \{1, 2, \ldots, n, \ldots, N\}$ is the set of players, $S = S_1 \times S_2 \times \cdots \times S_N$ is the set of profiles of pure strategies with S_n the finite set of pure strategies of player n, and $u = (u_1, u_2, \ldots, u_N)$ with $u_n : S \to \mathbb{R}$ the utility or payoff function of player n. We call the pair (N, S) the *game form*. Thus a game is a game form together with a payoff function.

Thus we have specified a set of players and numbered them 1 through N. Somewhat abusively we have also denoted this set by N. (This is a fairly common practice in mathematics and usually creates no confusion.) For each player we have specified a finite set of actions or strategies that the player could take, which we denote S_n. We have denoted the cartesian product of these sets by S. Thus a typical element of S is $s = (s_1, s_2, \ldots, s_N)$ where each s_n is a pure strategy of player n, that is, an element of S_n. We call such an s a pure strategy profile.

For each player n we have also specified a utility function $u_n : S \to \mathbb{R}$. We shall shortly define also randomised or mixed strategies, so that each player will form a probabilistic assessment over what the other players will do. Thus when a player chooses one of his own strategies he is choosing

a lottery over pure strategy profiles. So we are interpreting the utility function as a representation of the player's preferences over lotteries, that is, as a von Neumann-Morgenstern utility function.

DEFINITION 2.2. A *mixed strategy* of player n is a lottery over the pure strategies of player n. One of player n's mixed strategies is denoted σ_n and the set of all player n's mixed strategies is denoted Σ_n. Thus $\sigma_n = (\sigma_n(s_n^1), \sigma_n(s_n^2), \ldots, \sigma_n(s_n^{K_n})$ where K_n is the number of pure strategies of player n and $\sigma_n(s_n^i) \geq 0$ for $i = 1, 2, \ldots, K_n$ and $\sum i = 1^{K_n} \sigma_n(s_n^i) = 1$. The cartesian product $\Sigma = \Sigma_1 \times \Sigma_2 \times \cdots \times \Sigma_N$ is the set of all mixed strategy profiles.

The definition of a mixed strategy should by now be a fairly familiar kind of thing. We have met such things a number of times already. The interpretation of the concept is also an interesting question, though perhaps the details are better left for other places. The original works on game theory treated the mixed strategies as literally randomisations by the player in question, and in a number of places one finds discussions of whether and why players would actually randomise.

Some more recent works interpret the randomised strategies as uncertainty in the minds of the other players as to what the player in question will actually do. This interpretation seems to me a bit more satisfactory. In any case one can quite profitably use the techniques without worrying too much about the interpretation.

Perhaps more important for us at the moment as we start to learn game theory is the idea of extending the utility function of a player from that defined on the pure strategy profiles to that defined on mixed strategies. We shall continue to use the same symbol u_n to represent the expected utility of player n as a function of the mixed strategy profile $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_N)$. Intuitively $u_n(\sigma)$ is just the expected value of $u_n(s)$ when s is a random variable with distribution given by σ. Thus

$$u_n(\sigma) = \sum_{s_1 \in S_1} \cdots \sum_{s_N \in S_N} \sigma_1(s_1) \ldots \sigma_N(s_N) u_n(s_1, \ldots, s_N).$$

We can in a similar way define u_n on a more general profile where for some n we have mixed strategies and for others pure strategies.

1.2. Examples. Let's look now at some examples.

EXAMPLE 2.1 (The Prisoner's Dilemma). Consider the situation in which two criminals are apprehended as they are making off with their loot. The police clearly have enough evidence to convict them of possession of stolen property but don't actually have evidence that they actually committed the crime. So they question the criminals in separate cells and offer each the following deal: Confess to the crime. If the other doesn't confess then we shall prosecute him on the basis of your evidence and he'll go to jail for 10 years. In gratitude for your cooperation we shall let you go free. If the other also confesses then we don't actually need your evidence so we shall prosecute both of you, but since you cooperated with us we shall arrange with the judge that you only get 9 years. We are offering the other criminal the same deal. If neither of you confess then we shall be able to convict you for possession of stolen goods and you will each get one year.

How should we model such a situation? In actuality, of course, even in such a situation a criminal would have many options. He could make counter proposals; he could confess to many other crimes as well; he could claim that his partner had committed many other crimes. We might, however, learn something about what to expect in this situation by analysing a model in which each criminal's choices were limited to confessing or not confessing. We shall also take the actions of the police as given and part of the environment and consider only the criminals as players. We shall

assume that the criminals care only about how much time they themselves spend in jail and not about how much time the other spends, and moreover that their preferences are represented by a von Neumann-Morgenstern utility function that assigns utility 0 to getting 10 years, utility 1 to getting 9 years, utility 9 to getting 1 year and utility 10 to getting off free.

Thus we can model the situation as a game in which $N = \{1, 2\}$, $S_1 = S_2 = \{C, D\}$ (for (Confess, Don't Confess)) and $u_1(C, C) = 1$, $u_1(C, D) = 10$, $u_1(D, C) = 0$, $u_1(D, D) = 9$, and $u_2(C, C) = 1$, $u_2(C, D) = 0$, $u_2(D, C) = 10$, $u_1(D, D) = 9$.

Such a game is often represented as a labelled matrix as shown in Figure 1. Here player 1's strategies are listed vertically and player 2's horizontally (and hence they are sometimes referred to as the row player and the column player). Each cell in the matrix contains a pair x, y listing first player 1's payoff in that cell and then player 2's.

$1\backslash 2$	C	D
C	$1, 1$	$10, 0$
D	$0, 10$	$9, 9$

Figure 1

EXAMPLE 2.2 (Matching Pennies). This is a parlour game between two players in which each player chooses to simultaneously announce either "Heads" or "Tails" (perhaps by producing a coin with the appropriate face up). If the two announcements match then player 1 receives \$1 from player 2. If they don't match then player 2 receives \$1 from player 1.

We again assume that the players have only the two options stated. (They cannot refuse to play, go out and get drunk, or anything else equally enjoyable. They have to play the silly game.) Moreover, since the stakes are so low it seems a reasonable approximation to assume that their preferences are represented by a von Neumann-Morgenstern utility function that assigns utility -1 to paying a dollar and utility 1 to getting a dollar. (In fact, with only two possible outcomes this assumption of risk neutrality is unnecessary.)

Thus we can model the situation as a game in which $N = \{1, 2\}$, $S_1 = S_2 = \{H, T\}$ and $u_1(H, H) = 1$, $u_1(H, T) = -1$, $u_1(T, H) = -1$, $u_1(T, T) = 1$, and $u_2(H, H) = -1$, $u_2(H, T) = 1$, $u_2(T, H) = 1$, $u_1(T, T) = -1$. The game is represented by the labelled matrix shown in Figure 2.

$1\backslash 2$	H	T
H	$1,-1$	$-1,1$
T	$-1,1$	$1,-1$

Figure 2

1.3. Solution Concepts I: Pre-equilibrium Ideas. The central solution concept in noncooperative game theory is that of Nash equilibrium or strategic equilibrium. Before discussing the idea of equilibrium we shall look at a weaker solution concepts. One way of thinking of this concept is as the necessary implications of assuming that the players know the game, including the rationality and knowledge of the others.

There is quite a bit entailed in such an assumption. Suppose that some player,say player 1 knows some fact F. Now since we assume that the players know the knowledge of the other players the other players both know F and know that player 1 knows it. But then this is part of their knowledge and so they all know that they all know F. And they all know this. And so on. Such a situation was formally analysed in the context of game theory by Aumann (1976) who described F as being *common knowledge*.

Consider the problem of a player in some game. Except in the most trivial cases the set of strategies that he will be prepared to play will depend on his assessment of what the other players will do. However it *is* possible to say a little. If some strategy was strictly preferred by him to another strategy s whatever he thought the other players would do, then he surely would not play s. And this remains true if it was some lottery over his strategies that was strictly preferred to s. We call a strategy such as s a *strictly dominated strategy*.

Thus we have identified a set of strategies that we argue a rational player would not play. But since everything about the game, including

the rationality of the players, is assumed to be common knowledge no player should put positive weight, in his assessment of what the other players might do, on such a strategy. And we can again ask: Are there any strategies that are strictly dominated when we restrict attention to the assessments that put weight only on those strategies of the others that are not strictly dominated. If so, a rational player who knew the rationality of the others would surely not play such a strategy.

And we can continue for an arbitrary number of rounds. If there is ever a round in which we don't find any new strategies that will not be played by rational players commonly knowing the rationality of the others, we would never again "eliminate" a strategy. Thus, since we start with a finite number of strategies, the process must eventually terminate. We call the strategies that remain *iteratively undominated* or *correlatedly rationalizable*.

There is another related concept called *rationalizable strategies* that was introduced by Bernheim (1984) and Pearce (1984). That concept is both a little more complicated to define and, in my view, somewhat less well motivated so we won't go into it here.

1.4. Solution Concepts II: Equilibrium. In some games the iterative deletion of dominated strategies is reasonably powerful and may indeed let us say what will happen in the game. In other games it says little or nothing.

A more widely used concept is that of Nash equilibrium or strategic equilibrium, first defined by John Nash in the early 1950s. (See Nash (1950, 1951).) An equilibrium is a profile of mixed strategies, one for each player, with the property that if each player's uncertainty about what the others will do is represented by the profile of mixed strategies then his mixed strategy puts positive weight only on those pure strategies that give him his maximum expected utility. We can state this in a little more detail using the notation developed above.

DEFINITION 2.3. A *strategic equilibrium* (or *Nash equilibrium*) of a game (N, S, u) is a profile of mixed strategies $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_N)$ such that for each $n = 1, 2, \ldots, N$ for each s_n and t_n in S_n if $\sigma_n(s_n) > 0$ then

$$u_n(\sigma_1, \ldots, \sigma_{n-1}, s_n, \sigma_{n+1}, \ldots, \sigma_N) \geq u_n(\sigma_1, \ldots, \sigma_{n-1}, t_n, \sigma_{n+1}, \ldots, \sigma_N).$$

Remember how we extended the definition of u_n from the pure strategies to the mixed strategies at the beginning of this chapter.

We can relate this concept to that discussed in the previous section.

PROPOSITION 2.1. *Any strategic equilibrium profile consists of iteratively undominated strategies.*

Let's look back now to our two examples and calculate the equilibria. In both of these examples there is a unique equilibrium. This is not generally the case.

In the prisoner's dilemma the strategy of Don't Confess is dominated. That is, Confess is better whatever the other player is doing. Thus for each player Confess is the only undominated strategy and hence the only iteratively undominated strategy. Thus by the previous proposition (Confess,Confess) is the only equilibrium.

In matching pennies there are no dominated strategies. Examining the game we see there can be no equilibrium in which either player's assessment of the other's choice is a pure strategy. Let us suppose that suppose that player 1's assessment of player 2 was that player 2 would play H. Then player 1 would strictly prefer to play H. But then the definition of equilibrium would say that player 2 should put positive weight in his assessment of what player 1 would play only on H. And in this case player 2 would strictly prefer to play T, contradicting our supposition that player 1 assessed him as playing H. In fact we could have started this chain of argument by supposing that player 1's assessment of player 2 put weight strictly greater than a half on the fact that player 2 would play H. We could make a similar argument starting with the supposition that player 1's assessment of player 2 put weight strictly less than a half on the fact that player 2 would play H. Thus player 1's assessment of player 2's choices must be $(\frac{1}{2}, \frac{1}{2})$. And similarly player 2's assessment of player 1's choices must also be $(\frac{1}{2}, \frac{1}{2})$.

EXERCISE 2.1. Consider the variant of matching pennies given in Figure 3. Calculate the equilibrium for this game.

$1\backslash 2$	H	T
H	$1, -1$	$-2, 3$
T	$-1, 1$	$1, -1$

Figure 3

2. Extensive Form Games

There is another model of interactive decision situations that is, in some respects, perhaps a little more natural. In this model rather than listing all the plans that the decision makers might have and associating expected payoffs with each profile of plans one describes sequentially what each player might do and what results. The process is modelled as a multi-player decision tree. (You might possibly have come across the use of decision trees to describe decision problems facing single decision makers.)

We shall define such game trees or extensive form games in the next section. For now let us consider an example, that of Figure 4, that illustrates the essential ingredients.

The game starts at the bottom at Player 1's first decision node. The first node of the game is called the *initial node* or the *root*. Player 1 chooses whether to play T or B. If he chooses B he moves again and chooses between U and D. If he chooses B and then D the game ends. If he chooses either T or B and then U then player 2 gets a move and chooses either L or R. Player 2 might be at either of two nodes when she chooses. The dotted line between those nodes indicates that the are *in the same information set* and that Player 2 does not *observe* which of the nodes she is at when she moves. Traditionally information sets were indicated by enclosing the nodes of the information set in a dashed oval. The manner I have indicated is a newer notation and might have been introduced because it's a bit easier to generate on the computer, and looks a bit neater. (Anyway that's why I do it that way.)

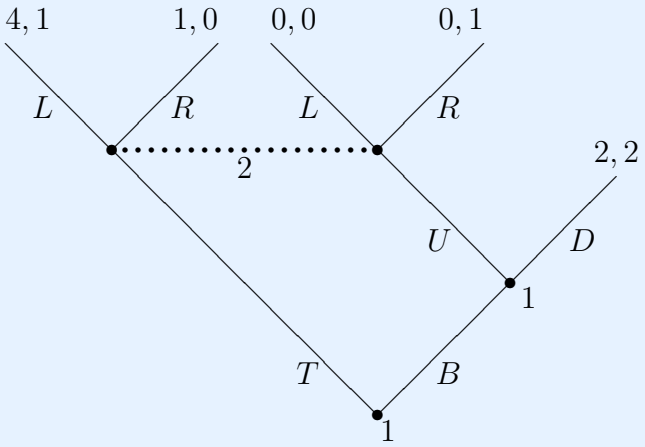

Figure 4

The payoffs given at the *terminal nodes* give the expected payoffs (first for Player 1 then for Player 2) that generate the von Neumann-Morgenstern utility function that represents the players' preferences over lotteries over the various outcomes or terminal nodes.

There are two further features that are not illustrated in this example. We often want to include in our model some extrinsic uncertainty, that is some random event not under the control of the players. We indicate this by allowing nodes to be owned by an artificial player that we call "Nature" and sometimes index as Player 0. Nature's moves are not labelled in the same

way as the moves of the strategic players. Rather we associate probabilities to each of nature's moves. This is shown in the game of Figure 5 in which the initial node is a move of nature. I shall indicate nodes where Nature moves by open circles and nodes where real players move by filled circles.

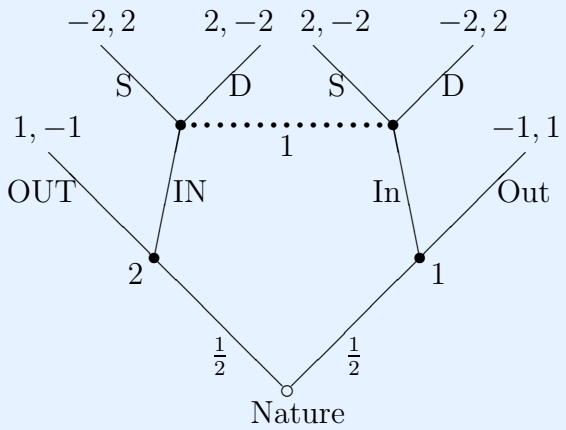

Figure 5

Like the Prisoner's Dilemma or Matching Pennies there is a bit of a story to go with this game. It's some kind of parlour game. The game has two players. We first assign a high card and a low card to players 1 and 2, each being equally likely to get the high card, and each seeing

the card he gets. The player receiving the low card then has the option of either continuing the game (by playing "In") or finishing the game (by playing "Out"), in which case he pays \$1 to the other player. If the player who received the low card continues the game then Player 1, moves again and decides whether to keep the card he has or to swap it with Player 2's. *However Player 1 does not observe which card he has, or what he did in his previous move, or even if he has moved previously.* This might seem a bit strange. One somewhat natural interpretation is that Player 1 is not a single person, but rather a team consisting of two people. In any case it is normal in defining extensive form games to allow such circumstances.

Games such as this are not, however, as well behaved as games in which such things do not happen. (We'll discuss this in a little more detail below.) If a player always remembers everything he knew and everything he did in the past we say that the player has *perfect recall*. If each player has perfect recall then we say that the extensive form game has perfect recall. The game of Figure 4 has perfect recall while the game of Figure 5 does not have prefect recall. In particular, Player 1 does not have perfect recall.

The extensive form given provides one way of modelling or viewing the strategic interaction. A somewhat more abstract and less detailed vision is provided by thinking of the players as formulating plans or strategies. One might argue (correctly, in my view) that since the player can when formulating his plan anticipate any contingencies that he might face nothing of essence is lost in doing this. We shall call such a plan a strategy. In the game of Figure 5 Player 2 has only one information set at which

she moves so her plan will simply say what she should do at that information set. Thus Player 2's strategy set is $S_2 = \{\text{IN}, \text{OUT}\}$. Player 1 on the other hand has two information sets at which he might move. Thus his plan must say what to do at each of his information sets. Let us list first what he will do at his first (singleton) information set and second what he will do at his second information set. His strategy set is $S_1 = \{(\text{In}, \text{S}), (\text{In}, \text{D}), (\text{Out}, \text{S}), (\text{Out}, \text{D})\}$.

Now, a strategy profile such as $((\text{In}, \text{S}), \text{IN})$ defines for us a lottery over the terminal nodes, and hence over profiles of payoffs. In this case it is $(-2, 2)$ with probability a half and $(2, -2)$ with probability a half. For the strategy profile $((\text{In}, \text{S}), \text{OUT})$ it would be $(1, -1)$ with probability a half and $(2, -2)$ with probability a half.

We can then calculate the expected payoff profile to each of the lotteries associated with strategy profiles (for the two given above this would be $(0, 0$ and $(1\frac{1}{2}, -1\frac{1}{2})$) and thus we have specified a normal form game. For this example the associated normal form game is given in Figure 5a.

$1\backslash 2$	IN	OUT
(In, S)	$0, 0$	$1\frac{1}{2}, -1\frac{1}{2}$
(In, D)	$0, 0$	$-\frac{1}{2}, \frac{1}{2}$
(Out, S)	$-1\frac{1}{2}, 1\frac{1}{2}$	$0, 0$
(Out, D)	$\frac{1}{2}, -\frac{1}{2}$	$0, 0$

Figure 5a

2.1. Definition. We now define formally the notions we discussed informally above.

DEFINITION 2.4. An *extensive form game* consists of

1. $N = \{1, 2, \ldots, N\}$ a set of players,
2. X a finite set of nodes,
3. $p : X \to X \cup \{\emptyset\}$ a function giving the immediate predecessor of each node. There is a single node x_0 for which $p(x) = \emptyset$. This is the initial node. We let $s(x) = p^{-1}(x) = \{y \in X \mid p(y) = x\}$ be the immediate successors of node x. We can now define the set of all predecessors of x to be those y's for which $y = p(p(p \ldots (x)))$ for some number of iterations of p and similarly the set of all successors of x. We require that for any x the set of all predecessors of x be disjoint from the set of all successors of x. (This is what we mean by the nodes forming a tree.) The set of terminal nodes T is the set of nodes that have no successors, that is those x for which $s(x) = \emptyset$. We call the nonterminal nodes the *decision nodes*.
4. A a set of actions and $\alpha : X \backslash \{x_0\} \to A$ a function that for any non-initial node gives the action taken at the preceding node that leads to that node. We require that if x and x' have the same predecessor and $x \neq x'$ then $\alpha(x) \neq \alpha(x')$. The set of choices available at the node x is $c(x) = \{a \in A \mid a = \alpha(x') \text{ for some } x' \in s(x)\}$.

5. $\mathcal{H}$ a collection of information sets, and $H : X \backslash T \to \mathcal{H}$ a function assigning each decision node x to an information set $H(x)$. We require that any two decision nodes in the same information set have the same available choices. That is if $H(x) = H(x')$ then $c(x) = c(x')$. We also require that any nodes in the same information set be neither predecessors nor successors of each other. Sometimes this requirement is not made part of the definition of a game but rather separated and used to distinguish *linear* games from *nonlinear* ones. (The linear ones are the ones that satisfy the requirement.)
6. $n : \mathcal{H} \to \{0, 1, \ldots, N\}$ a function assigning each information set to the player who moves at that information set or to Nature (Player 0). The collection of Player n's information sets is denoted $\mathcal{H}_n = \{H \in \mathcal{H} \mid n(H) = n\}$. We assume that each information set in $\mathcal{H}_0$ is a singleton, that is that it contains only a single node.
7. $\rho : \mathcal{H}_0 \times A \to [0, 1]$ a function that gives the probability of each of Nature's choices at the nodes at which Nature moves.
8. $u = (u_1, u_2, \ldots, u_N)$ a collection of payoff functions $u_n : T \to \mathbb{R}$ assigning an expected utility to each terminal node for each player n.

DEFINITION 2.5. Given an extensive form game we say that Player n in that game has perfect recall if whenever $H(x) = H(x') \in \mathcal{H}_n$ with x'' a predecessor of x and with $H(x'') \in \mathcal{H}_n$ and a'' the action at x'' on the path to x then there is $x''' \in H(x'')$ a predecessor to x' with action a'' the action at x''' on the path to x'. If each player in N has prefect recall we say the game has perfect recall.

You should go back to the games of Figure 4 and Figure 5 and see how this definition leads to the conclusion that the first is a game with perfect recall and the second is not. To understand the definition a little better let's look a little more closely at what it is saying. Since $H(x) = H(x')$ it means that the player observes the same situation at x and x'. Thus if he has perfect recall he should have the same experience at x and x'. Part of his experience at x was that he had been at the information set $H(x'')$ and made the choice a''. The definition is requiring that he also have had this experience at x'.

2.2. The Associated Normal Form. Just as in the previous section we formally defined the details of a game that we had earlier discussed informally here we shall formally define the process of associating a normal form game to a given extensive form game.

Recall that a normal form game has three components: a set of players, a strategy set for each player, and a utility function for each player giving that player's utility for each profile of strategies. The easiest part is defining the player set. It is the same as the player set of the given extensive form game. As I said above, a strategy for a player is a rule that tells the player what to do at each of his information sets.

DEFINITION 2.6. Given an extensive form game a strategy of Player n is a function $s_n : \mathcal{H}_n \to A$ with $s_n(H) \in c(H)$ for all H in $\mathcal{H}_n$.

So, we have now defined the second component of a normal form game. Now, a strategy profile specifies an action at each move by one of the "real" players and so defines for us a lottery over the terminal nodes. (It defines a lottery rather than simply a terminal node because we allow random moves by nature in our description of an extensive form game. In a game without moves by Nature a strategy profile would define a single terminal node.)

We then associate an expected payoff profile with the profile of strategies by taking the expected payoff to the terminal node under this lottery.

It might be a good idea to go back and look again at what we did in defining the normal form game associated with the extensive form game given in Figure 5.

2.3. Solution Concepts. Still to come.

3. Existence of Equilibrium

We shall examine a little informally in this section the question of the existence of equilibrium. Lets look first in some detail at an example. I shall include afterwards a discussion of the general result and a sketch of the proof. Remember, however, that we didn't do this in class and you are not required to know it for this course.

Let us consider an example (Figure 6). Player 1 chooses the row and player 2 (simultaneously) chooses the column. The resulting payoffs are indicated in the appropriate box of the matrix, with player 1's payoff appearing first.

	L	R
T	$2,0$	$0,1$
B	$0,1$	$1,0$

Figure 6

What probabilities could characterize a self-enforcing assessment? A (mixed) strategy for player 1 (that is, an assessment by 2 of how 1 would play) is a vector $(x, 1-x)$, where x lies between 0 and 1 and denotes the probability of playing T. Similarly, a strategy for 2 is a vector $(y, 1-y)$.

Now, given x, the payoff-maximizing value of y is indicated in Figure 6a, and given y the payoff-maximizing-value of x is indicated in Figure 6b. When the figures are combined as in Figure 6c, it is evident that the game possesses a single equilibrium, namely $x = \frac{1}{2}, y = \frac{1}{3}$. Thus in a self-enforcing assessment Player 1 must assign a probability of $\frac{1}{3}$ to 2's playing L, and player 2 must assign a probability of $\frac{1}{2}$ to Player 1's playing T.

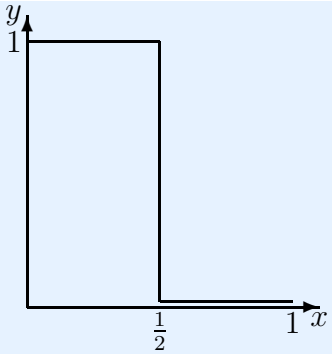

Figure 6a

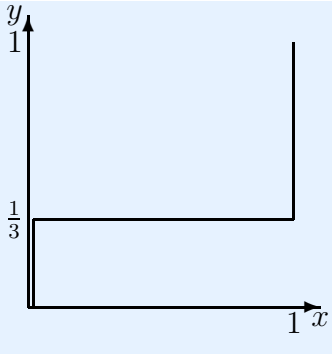

Figure 6b

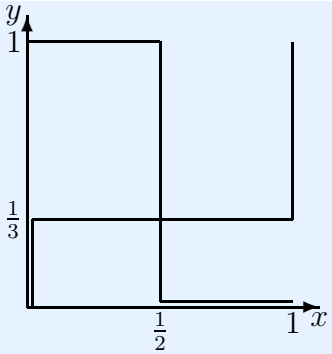

Figure 6c

The game of Figure 6 is an instance in which our notion of equilibrium completely pins down the solution. In general, we cannot expect such a sharp conclusion. Consider, for example, the game of Figure 7. There

are three equilibrium outcomes: (8,5), (7,6) and (6,3) (for the latter, the probability of T must lie between .5 and .6).

	L	C	R
T	$8,5$	$0,0$	$6,3$
B	$0,0$	$7,6$	$6,3$

Figure 7

In Figure 6 we see that while there is no pure strategy equilibrium there is however a mixed strategy equilibrium. The main result of non-cooperative game theory states that this is true quite generally.

THEOREM 2.1 (Nash 1950, 1951). *The mixed extension of every finite game has at least one strategic equilibrium.*

(A game is *finite* if the player set as well as the set of strategies available to each player is finite. Remember too that this proof is *not* required for this course.)

SKETCH OF PROOF. The proof may be sketched as follows. (It is a multi-dimensional version of Figure 6c.) Consider the set-valued mapping

(or correspondence) that maps each strategy profile, x, to all strategy profiles in which each player's component strategy is a *best response* to x (that is, maximises the player's payoff given that the others are adopting their components of x). If a strategy profile is contained in the set to which it is mapped (is a *fixed point*) then it is an equilibrium. This is so because a strategic equilibrium is, in effect, defined as a profile that is a best response to itself.

Thus the proof of existence of equilibrium amounts to a demonstration that the "best response correspondence" has a fixed point. The fixed-point theorem of Kakutani (1941) asserts the existence of a fixed point for every correspondence from a convex and compact subset of Euclidean space into itself, provided two conditions hold. One, the image of every point must be convex. And two, the graph of the correspondence (the set of pairs (x, y) where y is in the image of x) must be closed.

Now, in the mixed extension of a finite game, the strategy set of each player consists of all vectors (with as many components as there are pure strategies) of non-negative numbers that sum to 1; that is, it is a simplex. Thus the set of all strategy profiles is a product of simplices. In particular, it is a convex and compact subset of Euclidean space.

Given a particular choice of strategies by the other players, a player's best responses consist of all (mixed) strategies that put positive weight only on those pure strategies that yield the highest expected payoff among all the pure strategies. Thus the set of best responses is a subsimplex. In particular, it is convex.

Finally, note that the conditions that must be met for a given strategy to be a best response to a given profile are all weak polynomial inequalities, so the graph of the best response correspondence is closed.

Thus all the conditions of Kakutani's theorem hold, and this completes the proof of Nash's theorem. □

CHAPTER 3

Auctions

1. Introduction

Auctions are an extremely old economic institution. They were reported to have been used in ancient Babylonia where men bid for potential wives. Some have suggested that the biblical account of the sale of Joseph (the great grandson of Abraham) into slavery was an auction sale. Cassady (1967) contains a wealth of historical and recent examples.

In modern times auctions are used in a wide range of situations, both for rare and durable objects such as works of art with few substitutes and for objects such as a particular batch of agricultural produce that has many close substitutes. Auctions are used in situations in which the seller is in a strong position and when the seller is in a very weak position.

Widely used actions differ in a number of dimensions. There are both open or sequential auctions in which the bids are called out or or current bids announced, with all bidders hearing the announcements or bids; and sealed bid auctions in which bidders submit bids in a form that is not observable to the other bidders and the person auctioning the object selects one of the bids (or perhaps rejects all the bids).

In open auctions the most common form is the ascending bid auction or the English auction, in which bidders announce (or perhaps agree to) higher and higher bids until a stage is reached at which no one is willing to bid higher than the current highest bid. That bid then wins the auction (or perhaps all the bids are rejected). Another form of auction is the descending bid auction or the Dutch auction, in which the auctioneer starts by announcing a price higher than anyone would be willing to bid. He then

sequentially announces lower and lower prices until someone is willing to take that price (or reaches a price below which the seller is not willing to sell).

There are also two basic kinds of sealed bid auctions, the first price auction and the second price auction. In both the item is sold to the highest bidder. In the first price auction the highest bidder pays the price that he bid while in the second price auction he pays the price bid by the next highest bidder.

There are a number of good surveys of auction theory. One of the most readable is McAfee and McMillan (1987), although it does contain a little formal analysis a bit beyond the level that we shall do here. A slightly earlier book length survey is Engelbrecht-Wiggans, Shubik, and Stark (1983). The classic article in the literature is Vickrey (1961), which is worth taking a look at. (Vickrey was awarded the Nobel prize in 1996, at least in part for this work.)

The topic of auctions is currently a hot one because of the billions of dollars raised in the auctions of the airwave spectrum. A relatively nontechnical survey of the topic is McAfee and McMillan (1996). These auctions involve situations far more complicated than the simple ones we shall consider, and the theory is thus far more difficult and less satisfactorily worked out. This would be a challenging, but worthwhile topic for postgraduate research.

2. Types of Auction

There are many different arrangements that could be and are called auctions. Here we shall limit our discussion to a particularly simple case and discuss four different types of auction, two of which are very similar in their essential elements.

Let us consider a seller who has a single item to sell and that there is an identified set of potential buyers. We'll first divide auctions into those in which the potential buyers gather at some location (or perhaps some virtual location on the internet) so that they can observe bids being made and those in which the bidding is unobservable and conducted in a single round in which the potential bidders simultaneously submit their bids without the chance to observe what the other bidders are doing. In each case we shall divide the auctions into two further types.

In the sequential auctions we shall consider ascending-bid or English auction and the descending bid or Dutch auction. The English auction is far more widely used.

In the English auction the bids are successively raised until only one bidder is willing to bid at the current level. The bidding then stops and that bidder obtains the object for the last bid. The bids are either announced by the auctioneer with the bidders then indicating their willingness to accept that bid, or called out by the bidders themselves. It is not completely obvious in this case how we should model the situation. The actual sequential structure is potentially quite complicated, though we can argue (and shall) in some simple situations that the optimal behaviour is quite simple

to characterise. Nor is it always clear exactly how much information the bidders have. At a minimum they should all know the current bid. (This is what makes an English auction an English auction.)

In the Dutch auction the auctioneer initially calls out a high price and then successively lowers it until one bidder indicates that he will accept that price. In a Dutch auction the only information that bidders obtain is that the auction is over. The first time anyone observes anyone else bidding the game is over.

In sealed bid auctions (or tenders) all bidders submit a single bid, and do so simultaneously (i.e., without observing the bids of the others, or even if the others have bid). There are again two types of sealed bid auctions, the first price auction and the second price (or Vickrey) auction. In both types of sealed bid auction the highest bidder obtains the object. In the first price auction he pays the price he bid. In the second price auction he pays the highest price bid by anyone *else*.

Let us now consider the Dutch auction and the first price sealed bid auctions. What strategies are available to the bidders? Well in each case a bidder obtains no usable information as the game progresses. In the Dutch auction he is told a number of times that the game is still progressing, but he could anticipate in advance that he might have had this information. At some stage he will either choose to bid, and thus end the game, or be told that someone else has ended the game. Essentially all he can do is decide at what level he will be prepared to bid, given that no one has preempted him by bidding at a higher level. Thus he will choose a bid. The auction

will be won by the bidder who chooses the highest bid and that bidder will pay the price he bid.

Notice that this description of strategies, and the rule that tells us the outcome as a function of the strategies is exactly what happens in the first price sealed bid auction. Each bidder chooses a bid and the auction will be won by the bidder who chooses the highest bid and that bidder will pay the price he bid. The Dutch auction and the first price sealed bid auction are *strategically equivalent*.

Recall what we said earlier about the English auction. It isn't clear exactly what information the bidders have, or what they can do. They might or might not be able to perfectly observe who is making what bids. Let us consider the following simple (and unrealistic) auction form. Each bidder sits at a table with a button. The auctioneer starts at a very low price (perhaps zero) and successively increases the price. At each level if the bidder is willing to bid at least that price he does nothing. If he is not he pushes the button. Having once indicated that he is not willing to bid the bidder can never re-enter the bidding. As soon as everyone but one bidder has dropped out the auction ends and the single bidder left wins the object and pays the price at which the second last bidder dropped out.

As with the Dutch auction above this is strategically equivalent to a sealed bid auction, but this time to the second price sealed bid auction. In some circumstances other models of the English auction will have the same equilibria as this model, however, they offer a richer set of strategies.

Potentially the bids could depend on any information the bidder has obtained. For example he might follow the (partially described) strategy of bidding up to \$10 as long as there were at least 4 bidders still active at \$5. Such a strategy does not have a counterpart in the sealed bid auction.

3. Analysis of the Auctions I: Private Values

We considered in the previous section the structure of the bidding situation, what the bidders could do and what the outcome would be given what each bidder chose to do. In this section we shall consider the bidders preferences and we shall make very strong and simple assumptions about these preferences. This is not meant to be particularly realistic. Rather it will give us some understanding of what is involved in solving such models, even in the simplest cases, and will give us a benchmark against which to measure other situations.

As well as saying what the preferences of the bidders are, at least for some of the auction forms, we also need to consider what the bidders know about the preferences of the other bidders, what they know about what the others know about their preferences, and so on.

We suppose that there are N potential bidders and that the preferences of the bidders are determined as follows. Bidder n's preferences are described by a utility function of the form

$$U_n = \begin{cases} V_n - P \text{ if bidder } n \text{ obtains the object and pays an amount } P \\ -P \text{ if bidder } n \text{ does not obtain the object and pays an amount } P. \end{cases}$$

Alternately, we can write this as

$$U_n(W, P_n) = 1_{\{W=n\}} V_n - P_n$$

where W is the bidder who obtains the object, $1_{\{W=n\}} = 1$ if $W = n$ and 0 otherwise, and P_n is the amount paid by bidder n. Notice that we define

the preferences over more situations than can obtain given the auctions we have so far described. In particular we allow for cases in which the bidders may pay some amount, even if they do not obtain the object.

This part of the description of preferences we shall maintain for most of these notes. Its content is basically that the bidders are risk neutral.

We now look at how the value V_n is determined and what the bidders know about the vector $V = (V_1, V_2, \ldots, V_N)$. We could, without making the problem too difficult be a bit more general, but we shall make the following very simple assumption.

Each V_n is randomly determined according to a uniform distribution on $[0, \bar{V}]$ with V_n independent of V_m whenever $n \neq m$. Each bidder knows this and bidder n knows the realised value of V_n, but nothing further about the rest of V.

Let us now consider the two essentially different forms of auction that we described in the previous section. We start with the second price auction, since it is easier.

3.1. Second Price Auction. Each bidder knows his own valuation. Thus anything he finds out, or even anything he could find out, about the bids of the others does not affect whether or not he would be willing to accept the object for a particular price. Whether or not he wants to be the winner of the auction depends only on the price he will have to pay if he is the winner. Now, when does bidder n want to be the winner? Clearly whenever he will have to pay less than his valuation V_n. And he will want to be one of the losers if he would have to pay more than V_n. But in the second price auction the price he would have to pay is determined only by the bids of the others. If the highest bid of any other bidder is \$7 then bidder n will pay \$7 if he bids higher than \$7 and will not win if he bids lower than \$7. His own bid only determines whether or not he will be the winner. (And, of course, might determine the price someone else pays if he is not the winner.)

Suppose that the bidder knew exactly what the other bidders had bid. If the highest bid of the others is less than his valuation then he will want to bid strictly more than this bid and he will be indifferent among all bids that are strictly more than the highest other bid. In particular bidding his true valuation is one of his best choices.

On the other hand if the highest bid of the others is more than his valuation then he will want to bid strictly less than this bid and he will be indifferent among all bids that are strictly less than the highest other bid. Again, bidding his true valuation is one of his best choices.

And if the highest bid of the others is exactly equal to his valuation then he will be indifferent among all bids. Once again, bidding his true valuation is one of his best choices.

Thus bidding his true valuation is always one of his best choices and any other bidding strategy will, in some circumstances, be strictly worse that bidding his true valuation.

We've been able to work out the equilibrium strategies in this case without even doing any maths. (I'm sure that this made you all happy.) Let's now try to work out the expected price to the seller. We've seen that each bidder will bid his true value. And the object is sold for the second highest bid. Thus we have to work out the expected value of the second highest bid.

This isn't too hard but we do need a bit of notation, and we need to persevere with a bit of calculation. The distribution of each bidder's value, V_n is the same. Remember that we assumed the V_n was distributed uniformly on $[0, \bar{V}]$. Thus

$$\Pr(V_n \leq x) = F(x) = \frac{x}{\bar{V}}$$

for $0 \leq x \leq V_n$. Let's introduce the following notation: $V_{(1)}$ is the highest of the V_n's, $V_{(2)}$ is the second highest, and so on. What we want is the expectation of $V_{(2)}$. Let's calculate the distribution function of $V_{(2)}$. What we need to calculate is $\Pr(V_{(2)} \leq x)$ for an arbitrary x. We divide the event that $V_{(2)} \leq x$ into $N + 1$ mutually exclusive and completely exhaustive

subsets (that is, if $V_{(2)} \leq x$ is true then exactly one of the subevents is true). One way that it could be that $V_{(2)} \leq x$ is that V_1 is greater than x and the remainder of the V_n are less than x. Another way is that V_2 is greater than x and the remainder of the V_n are less than x, and so on. Finally it might be that all of the V_n are less than x. The probabilities of each of these subevents are easy to calculate and since they are mutually exclusive the probability of the event that $V_{(2)} \leq x$ is simply the sum of these probabilities.

The probability that V_1 is greater than x is $1-F(x)$ and the probability that V_n is less than x is $F(x)$. And since the V_n's are all independent the probability that V_1 is greater than x and the remainder of the V_n are less than x is

$$\left(F(x)\right)^{N-1}\left(1-F(x)\right),$$

and similarly for each of the first N subevents. Finally, the probability that all of the V_n are less than x is $\left(F(x)\right)^N$. Putting all this together we obtain

$$\begin{aligned}\Pr(V_{(2)} \leq x) &= N\left(F(x)\right)^{N-1}\left(1-F(x)\right)+\left(F(x)\right)^N \\ &= \left(F(x)\right)^{N-1}\left(N\left(1-F(x)\right)+F(x)\right).\end{aligned}$$

And, recalling the formula for $F(x)$ this is equal to

$$N\left(\frac{x}{\bar{V}}\right)^{N-1}-(N-1)\left(\frac{x}{\bar{V}}\right)^N.$$

Let's call this distribution G.

To calculate the expected value of $V_{(2)}$ we need to integrate $V_{(2)}$ times g the density of $V_{(2)}$ where the density function g is simply the derivative of G. The details of this fairly easy calculation are left to you to do. The answer is that

$$g(x) = N(N-1)\frac{x^{N-2}}{\bar{V}^{N-1}}\left(1-\frac{x}{\bar{V}}\right).$$

And so

$$\begin{aligned}
\mathrm{E}\left(V_{(2)}\right) &= \int_0^{\bar{V}} x g(x)\mathrm{d}x \\
&= \int_0^{\bar{V}} N(N-1)\frac{x^{N-1}}{\bar{V}^{N-1}}\left(1-\frac{x}{\bar{V}}\right)\mathrm{d}x \\
&= \left(\frac{N-1}{N+1}\right)\bar{V}.
\end{aligned}$$

3.2. First Price Auction. Here we maintain the assumptions about preferences but consider the first price auction. In the second price auction, once we understood the idea, it was very easy to see the optimal strategies for the bidders. We did not have to exploit the idea of equilibrium; the bidders each had a strategy that was good for them, whatever the others bidders did. Such is not the case in the first price auction. Here each bidder really must consider what the other bidders will do.

Recall from the part of the course about game theory the definition of Nash equilibrium. A profile of strategies is a Nash equilibrium if the strategy of each player is at least as good as any other strategy of that player, *given the strategies of the others.* Now, a strategy of bidder n describes what bidder n will bid for any of the possible values V_n. That is, it is a function from the possible values, $[0, \bar{V}]$ to possible bids. (We haven't been too specific about the possible bids. Lets assume here that a bidder could bid any nonnegative amount.) That is, bidder n strategy is a function $B_n : [0, \bar{V}] \to \mathbb{R}_+$. It possible, but a bit more difficult, to completely solve for the equilibrium, without making further assumptions. However we'll immediately *assume* that

1. each bidder uses the same bidding strategy B, and
2. B is a strictly increasing function of the valuation V_n.

Now, since B is strictly increasing we can invert it. That is, B tells us, for any V_n what the bid of bidder n will be. The inverse function B^{-1} tells

us for any bid of bidder n what the valuation of bidder n is. Suppose we want to know the probability that bidder n will bid lower than \$7. This is precisely the probability that the valuation of bidder n is lower than the valuation that would have led him to bid \$7. That is,

$$\Pr(B_n \leq b) = \Pr(V_n \leq B^{-1}(b)) = F(B^{-1}(b)) = \frac{B^{-1}(b)}{\bar{V}}.$$

In fact, this *is* a probability in which any other bidder, say bidder 1, would be interested. It tells bidder 1 if he makes the bid b the probability that bidder n will bid less than b. (Since we have a continuous distribution and assume that B is strictly increasing the probability of a bid being exactly b is zero and so we ignore the possibility of ties.) Since all the other bidders are assumed to use the function B and their valuations all have the same distribution and are independent the probability that they *all* bid less than b is

$$\Pr(B_n \leq b \text{ for all } n = 2, 3, \ldots, N) = \left(\frac{B^{-1}(b)}{\bar{V}}\right)^{N-1}.$$

This is the probability that bidder 1, bidding b, will win the bid.

Now, what is bidder 1's expected payoff from bidding b. Well, if he wins the bid he gets $V_1 - b$ and if he loses the bid he gets zero. We've just seen what his probability of winning is. So his expected payoff is

$$(V_1 - b)\left(\frac{B^{-1}(b)}{\bar{V}}\right)^{N-1}. \tag{2}$$

Observe that his expected payoff is the product of two values. The first is his payoff if he wins and the second his probability of winning. Increasing his bid b lowers the first value and raises the second. Thus when he optimally chooses b any gain in one of the values from changing b is exactly offset by a loss in the other value. Or differentiating 2 and setting the result to zero we obtain that

$$(3)\quad \left(\frac{B^{-1}(b)}{\bar{V}}\right)^{N-1} = (V_1 - b)\,(N-1)\left(\frac{B^{-1}(b)}{\bar{V}}\right)^{N-2}\left(\frac{1}{\bar{V}}\right)\left(\frac{\mathrm{d}B^{-1}(b)}{\mathrm{d}b}\right)$$

for b equal to the optimal bid for a bidder with value V_1.

Now $\mathrm{d}B^{-1}(b)/\mathrm{d}b$ is the slope of the inverse of the bidding function at the bid b. That is, it's the inverse of (that is, one over) the slope of the bidding function B at the value $B^{-1}(b)$. Thus, simplifying 3 we obtain

$$(4)\qquad B^{-1}(b)\frac{\mathrm{d}B}{\mathrm{d}V}\left(B^{-1}(b)\right) = (V_1 - b)(N-1).$$

Since we have assumed that all the bidders are using the same bidding function B and that b is the optimal bid for a bidder with value V_1 we also have that $B^{-1}(b) = V_1$ and $b = B(V_1)$. Thus 4 becomes (writing V instead of V_1 since this is now true for all bidders)

$$(5)\qquad V\frac{\mathrm{d}B}{\mathrm{d}V}(V) = (V - B(V))(N-1).$$

or

$$\frac{\mathrm{d}B}{\mathrm{d}V}(V)V = \left(1 - \frac{B(V)}{V}\right)(N-1). \tag{6}$$

It is also clear that the a bidder with zero valuation will bid zero, that is $B(0) = 0$. Thus we have a first order ordinary differential equation 6 and one boundary condition. This can be (easily) solved. Solving differential equations is a bit beyond the scope of this course, so I'll just state the solution. You should confirm that it does indeed satisfy the differential equation and the boundary condition.

$$B(V) = \left(\frac{N-1}{N}\right)V.$$

This has a clear interpretation. A bidder's bid is related to his valuation. However, unlike the second price auction, he does not bid his true value. In the first price auction this would ensure that he never obtained a positive payoff. Rather he bids a little less than his true valuation. In the case we consider he bids a fraction $(N-1)/N$ of his true value. As we might expect we see that as the number of bidders increases he bids more aggressively.

It remains to calculate the expected price. After all this work, you'll be happy to hear that this bit is a little easier than in the second price auction. In the first price auction the price is equal to the winning bid. The highest possible bid is $(N-1)\bar{V}/N$. For any x between zero and this value the probability of the winning bid being less than this x is the probability of

all bids being less than x which, with the bidding function we found is the probability that all values are less than $Nx/(N-1)$. That is,

$$\begin{aligned}\Pr(\text{Winning Bid} \leq x) &= \Pr\left(\text{All Values} \leq \frac{Nx}{N-1}\right) \\ &= \left(\frac{Nx}{(N-1)\bar{V}}\right)^N.\end{aligned}$$

Let's again call this distribution G, and find the density function g by differentiating G.

$$g(x) = \left(\frac{N}{(N-1)\bar{V}}\right)^N Nx^{N-1}.$$

And so

$$\begin{aligned}\mathrm{E}\,(\text{Winning Bid}) &= \int_0^{\frac{(N-1)\bar{V}}{N}} xg(x)\mathrm{d}x \\ &= \int_0^{\frac{(N-1)\bar{V}}{N}} \left(\frac{N}{(N-1)\bar{V}}\right)^N Nx^N \mathrm{d}x \\ &= \left(\frac{N}{(N-1)\bar{V}}\right)^N \left(\frac{N}{N+1}\right)\left(\frac{N-1}{N}\bar{V}\right)^{N+1} \\ &= \left(\frac{N-1}{N+1}\right)\bar{V}.\end{aligned}$$

As we claimed earlier, the expected price is exactly the same as in the second price auction.

4. Bibliography

Robert J. Aumann (1976): "Agreeing to Disagree," *The Annals of Statistics*, 4, 1236–1239. 37

B. Douglas Bernheim (1984): "Rationalizable Strategic Behavior," *Econometrica*, 52, 1007–1028. 38

Jr. Cassady, Ralph (1967): *Auctions and Auctioneering*, University of California Press, Berkeley. 60

Richard Engelbrecht-Wiggans, Martin Shubik, and Robert M. Stark, editors (1983): *Auctions, Bidding, and Contracting: Uses and Theory*, New York University Press, New York. 61

Shizuo Kakutani (1941): "A Generalization of Brouwer's Fixed-Point Theorem," *Duke Mathematics Journal*, 8, 457–459. 57

R. Preston McAfee and John McMillan (1987): "Auctions and Bidding," *Journal of Economic Literature*, 25, 699–738. 61

R. Preston McAfee and John McMillan (1996): "Analyzing the Airwaves Auction," *Journal of Economic Perspectives*, 10, 159–175. 61

John Nash (1950): "Equilibrium Points in N-Person Games," *Proceedings of the National Academy of Sciences*, 36, 48–49. 39

John Nash (1951): "Non-Cooperative Games," *Annals of Mathematics*, 54, 286–295. 39

David G. Pearce (1984): "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52, 1029–1050. 38

William Vickrey (1961): "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16, 8–37. 61